ELSEVIER

Journal of Economic Dynamics and Control
21 (1997) 1377–1403

JOURNAL OF
Economic
Dynamics
& Control

Strategic asset allocation

Michael J. Brennan[a, b, *], Eduardo S. Schwartz[c], Ronald Lagnado[d]

[a]*Department of Banking and Finance, University of California, Los Angeles, USA*
[b]*Department of Finance, London Business School, London, UK*
[c]*Department of Real Estate and Finance, University of California, Los Angeles, USA*
[d]*CATS*

Abstract

This paper analyzes the portfolio problem of an investor who can invest in bonds, stock, and cash when there is time variation in expected returns on the asset classes. The time variation is assumed to be driven by three state variables, the short-term interest rate, the rate on long-term bonds, and the dividend yield on a stock portfolio, which are all assumed to follow a joint Markov process. The process is estimated from empirical data and the investor's optimal control problem is solved numerically for the resulting parameter values. The optimal portfolio proportions of an investor with a long horizon are compared with those of an investor with a short horizon such as is typically assumed in 'tactical asset allocation' models: they are found to be significantly different. Out of sample simulation results provide encouraging evidence that the predictability of asset returns is sufficient for strategies that take it into account to yield significant improvements in portfolio returns.

Keywords: Intertemporal portfolio theory; Stochastic control theory; Investment policy
JEL classification: G0, C6, D9

1. Introduction

One of the earliest applications of the portfolio theory developed by Markowitz was to 'tactical asset allocation', the systematic allocation of investment portfolios across broad asset classes such as bonds, stock and cash. Tactical asset allocation strategies have gained greatly in popularity in the wake

*Correspondence address: Michael J. Brennan, Anderson School of Management, UCLA, Los Angeles, CA 90095-1481, USA. Telephone and fax: (310) 206-8419; e-mail: michael.brennan@agsm.ucla.edu.

0165-1889/97/$17.00

of the Stock Market Crash of 1987, since those following the strategy managed largely to avoid an overcommitment to equities immediately before the Crash.

Tactical asset allocation is essentially a single period or myopic strategy; it assumes that the decision maker has a (mean–variance) criterion defined over the one period rate of return on the portfolio. This gives rise to two difficulties. First, the expected rates of return that are inputs to the model are typically not one-period expected rates of return, but rather estimated internal rates of return over long holding periods. For example, the expected return on bonds is usually proxied by the yield to maturity on a long-term bond, while the expected return on stocks is the constant discount rate implied by a dividend discount model. The implicit assumption in using these measures of expected return is that the one-period expected return is proportional to the estimated long-run internal rate of return – insofar as this assumption is not satisfied, the resulting portfolio will be biased towards one or another of the asset classes. Therefore, it would seem preferable to relate expected one-period rates of return to these proxies, for example by regression analysis, and to use the predicted one-period rate of return from the regression in the portfolio problem.

The second difficulty with tactical asset allocation is potentially more fundamental, in that it concerns the objective function. A *sine qua non* of tactical asset allocation is time variation or predictability in expected asset returns – a market in which asset returns conform to the random walk hypothesis would imply constant portfolio proportions, at least for iso-elastic utility functions, or when the mean variance criterion is defined over the rate of return on the portfolio. There is now widespread evidence of predictability in asset returns,[1] and Mossin (1968) has shown that a single period or myopic objective function of the type that underlies tactical asset allocation is appropriate only if the investor has a logarithmic utility function. For general (non-log) utility functions the investor will be concerned about hedging against shifts in the future investment opportunity set (changes in expected returns or covariances) – for an investor with a long horizon, a drop in interest rates may be as important for his future welfare as a substantial reduction in his current wealth. Similar considerations apply to institutional investors such as pension funds, depending on the precise specification of their objective function.[2] Merton (1971, 1990) has considered the problem of an investor planning his lifetime consumption and portfolio strategy.

[1]See for example Keim and Stambaugh (1986), Fama and French (1988), Campbell and Shiller (1988), Cochrane (1990), and Beckaert and Hodrick (1992).

[2]Merton (1990) has made the interesting point to us that a university endowment fund may wish to invest in property around the university in order to hedge against increases in property values that would inhibit their ability to expand or to attract faculty on account of housing costs in the locale.

Except for special cases[3] it is not possible to obtain closed-form solutions for the optimal strategies and, until recently, lack of computing power has made it impracticable to solve realistic problems with time varying expected returns.

In this paper we formulate and solve the portfolio problem of an investor who has a long-term horizon, with utility defined over wealth at the end of the horizon,[4] when the structure of expected returns may be described by a small number of state variables which follow a joint Markov process. The investor is assumed to invest in three assets – an instantaneously riskless security, 'cash', a long-term (consol) bond, and an equity portfolio. The variables that predict expected returns on these assets are the instantaneously riskless interest rate (the 'short rate'), the yield on the consol bond (the 'long rate'), and the dividend yield on the equity portfolio (the 'dividend yield'). While the model as formulated here assumes that the investor has no liabilities, it is relatively straightforward to generalize it to allow for liabilities whose expected rate of increase depends on the levels of the state variables – this will be important if we are to apply the model to the planning problem of a pension fund. Finally, we ignore here any problems of inflation; again, for practical implementation it will be important to incorporate stochastic inflation in the analysis. Our objective is to compare the portfolio strategies implied by a myopic objective function which ignores changes in the future investment opportunity set, with those implied by optimal behavior for an investor with a long–term horizon.

The stochastic optimal control problem that we formulate, following Merton (1971), may be contrasted with the stochastic programming approach which has been followed by several recent authors.[5] The essential difference between the two approaches is the way in which the uncertainty in the environment is modelled. The stochastic optimal control problem captures uncertainty by allowing for a continuum of states[6] which can be described at a given point in time by a small number of state variables that follow a joint Markov process; the size of the stochastic optimal control problem grows exponentially with the number of state variables, which limits the applicability of the approach to situations in which it is reasonable to model the state of the world by a relatively small number of state variables. In the application described here, the state of the world depends only on the investment opportunities which are captured by three state variables: two interest rates and the dividend yield on an equity portfolio. Given this, a problem with 480 time steps over 20 years can be solved in a matter of hours on a Pentium PC. However, such a parsimonious state

[3]See Merton (1990, Chapter 5).

[4]It is a simple extension to allow for utility from intermediate 'consumption' withdrawals.

[5]See for example Carino *et al.* (1994), and Mulvey and Vladimirou (1992).

[6]In practice, it is usually necessary to discretize the state space for computational purposes.

description would be inadequate, for example, if transactions costs were important,[7] or if the agent had non-tradable stochastic assets or liabilities, for then the exact composition of the asset or liability portfolio at each point in time would become relevant, and this would vastly expand the state space. On the other hand, the approach does allow consideration of a large number of investment assets, for the problem grows only linearly in the number of assets, and portfolio constraints are easy to implement provided that they depend on current market values and not on historical book values, for the book values would then become part of the state vector, again vastly expanding the state space.

Stochastic programming models capture uncertainty by a branching event tree. Each node of the tree represents a joint outcome of all the random variables at that decision stage, and corresponds to a particular realization of the state variables at a point in time in the stochastic optimal control approach. The major advantage of stochastic programming is that it can easily accommodate a large number of random variables at each node, and thus permits a very rich description of the state of the world, which may include the book value of all assets and liabilities (which may be important for regulatory reasons or on account of a capital gains tax), the age and sex composition of the workforce (which may be relevant for a pension fund), as well as market values of each asset and liability class which may be important if transaction costs are significant. Each path through the event tree represents a 'scenario'. The total number of scenarios depends mutiplicatively on the number of branches from each node and the number of time steps or decision nodes. For example, ten time steps and three branches from each node implies 59,049 different scenarios; while problems of this order of magnitude can now be solved,[8] they are highly computer intensive. Moreover, three branches (or even ten) emanating from a single node can provide only a very limited description of the uncertainty facing the decision maker over the next decision interval.[9] As Mulvey and Vladimirou (1992, p. 1661) point out, 'Statistical methods ... can limit the required number of scenarios to properly capture uncertainty and maintain computational tractability ... These issues are the subject of active research'. Thus stochastic optimal control and stochastic programming approaches can be viewed as complementary rather than competing. The former approach has significant computational advantages where the problem under consideration can be

[7]We do not regard transactions costs as important for trading between major asset classes in modern capital markets, for such portfolio adjustments can be effected by trading in futures contracts at minimal cost.

[8]Jessup *et al.* (1994) solve a problem with 98,000 scenarios in 23 seconds on the Intel iPSC/860 Connection Machine CM-5.

[9]For example, in the absence of transaction costs or constraints on portfolio holdings the optimization problem at a node with n branches can determine at most the values of the positions in n different assets.

adequately represented in terms of a small number of state variables that follow a joint Markov process. The latter is the only alternative where the number of state variables is large, and it has been suggested to us[10] that with developments in interior point algorithms, and with the potential use of parallel and distributed computations, stochastic programs with thousands or millions of scenarios become solvable.[11] Significant practical or commercial applications of stochastic programming include Carino et al. (1994) who describe a case study of a Japanese insurance company with 256 scenarios, with eight branches from the initial node and two or three thereafter, Mulvey (1994) who analyzes an asset–liability problem, Mulvey and Vladimirou (1992) who analyze a variety of financial planning problems with up to 8 periods and 100 scenarios, and Zenios (1993) who develops a model for the management of a portfolio of mortgage-backed securities. We are aware of no practical or commercial applications of stochastic optimal control theory in asset management.

In Section 1 we illustrate the essential difference between myopic portfolio selection and dynamic ('strategic') strategies, using a simple two-period example in which wealth may be invested in short- or long-term bonds. In Section 2 we present the optimal control problem. Section 3 specifies the stochastic process for the state variables and presents empirical estimates of the parameters. Section 4 discusses the solution of the control problem and Section 5 discusses the differences between the portfolio composition under an optimal policy with that under a myopic strategy.

2. A simple example

Consider an investor with initial wealth W who is interested in maximizing the expected utility of wealth at the end of a two period horizon. Assume that his utility function is of the iso-elastic family:

$$U(W) = \frac{1}{\gamma} W^{\gamma}. \tag{1}$$

At time t ($t = 0, 1, 2$) the investor's expected utility under the optimal policy will depend on both his current wealth, W, and the investment opportunities he faces, which we assume can be represented by a vector $\boldsymbol{Y}$, $V(W, \boldsymbol{Y}, t)$. We assume for illustrative purposes that the investor's investment opportunities are limited to a one-period bond and a two-period bond, and that the *pure expectations*

[10] By a referee of this journal

[11] See Mulvey and Ruszczynski (1995), and Jessup et al. (1994).

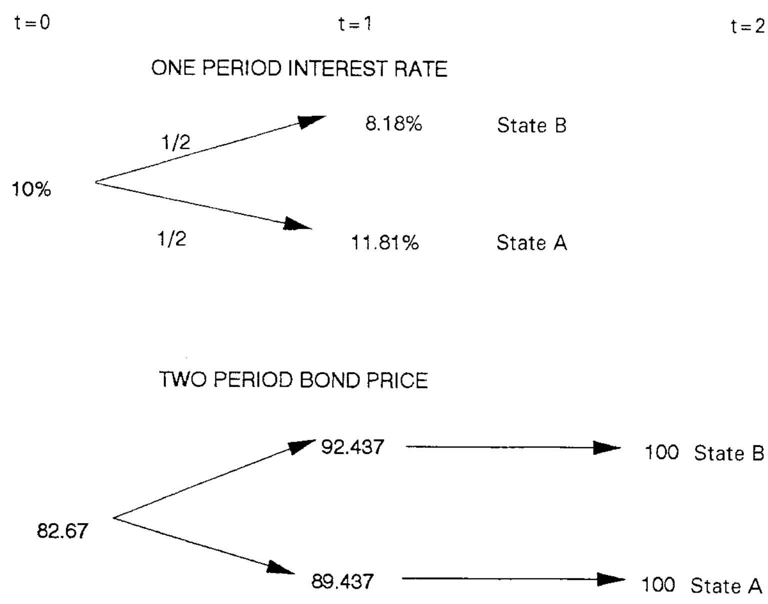

Fig. 1. Binomial model of bond pricing under the pure expectations hypothesis.

hypothesis holds so that the expected return on the one and two-period bonds are the same. The one-period interest rate is assumed to follow a binomial process with binomial probability of 1/2 as shown in Fig. 1. The two-period bond price at $t = 1$ is obtained by discounting the final payoff at the relevant one period rate; the price at $t = 0$ is obtained by discounting the expected value of the bond at $t = 1$ by the riskless rate of 10% at $t = 0$.[12] We have not considered a two-period bond in the second investment period, since a risk averse investor with a single-period horizon would never invest in an asset whose return is risky, if its expected return is equal to the riskless rate.[13]

First, note that a myopic risk averse investor will not invest anything in the two-period bond in the first period, since its return is risky and its expected return is only the same as the riskless return available on the one-period bond.

[12]It is at this stage that we are using the expectations hypothesis.

[13]See Arrow (1971).

Table 1
Optimal allocation to two-period bond (x) as a function of the coefficient of relative risk aversion (γ)

γ	0.9	0.5	0.0	− 0.5	− 0.9	− 2.0
x	− 8.9	− 1.0	0.0	0.3	0.4	0.7

However, an investor who behaves optimally will recognize that the long-term bond has a low return in State A, but that this is compensated for by the higher reinvestment rate in State A, and conversely for State B. Thus the two-period bond may allow the investor to hedge against changes in the future investment opportunity set as represented here by the one-period interest rate in the second period. As a result, the investor may find it advantageous to take a position in the two period bond in the first period. Since the investor invests only in the one period bond in the second period, his final wealth in the two states may be written as

$$W_A = 1.1181 W_0 \{1.1 + x(0.0818 - 0.1)\},$$

$$W_B = 1.0818 W_0 \{1.1 + x(0.1181 - 0.1)\},$$

where x is the fraction of wealth allocated to the two-period bond in the first period. The investor chooses x to maximize

$$V(W_0, 10\%, 0) = 0.5 W_A^{\gamma} + 0.5 W_B^{\gamma}. \tag{2}$$

The optimal values of x for different values of the risk aversion parameter, γ, are shown in Table 1. It is seen that for $\gamma > 0$ it is optimal for the investor to short the two-period bond, while for $\gamma < 0$, it is optimal for the investor to take a long position. Only for $\gamma = 0$, which corresponds to the log function, is it optimal for the investor to take no position – for this utility function the investor behaves myopically and, as we noted, a myopic risk averse investor will not invest in a risky asset unless it offers a positive risk premium. Thus, this example illustrates the main point of the paper, that whereas a myopic or single-period investor will treat a long-term bond as simply another risky asset, an investor with a long-term horizon will recognize also the ability of the bond to hedge against future changes in the investment opportunity set. In more complex settings other risky assets may have similar hedging characteristics.

3. The optimal control problem

In selecting state variables to represent the opportunity set of the investor we were guided by the need for parsimony, since the size of the control problem grows geometrically with the number of state variables. The short term interest

rate, r, was selected as a state variable, both because it is the return on an asset 1
class, and because there is extensive evidence that the level of the short rate predicts the expected return on common stock.[14] The second most powerful predictor of stock returns is the dividend yield on common stocks,[15] δ, and therefore this was included as the second state variable. Finally, we included the yield on a consol bond, l, as the third state variable, because of prior evidence that expected changes in the short rate are related to the current value of the long rate.[16] A fourth possible candidate is the junk bond yield spread, because there is evidence that this has predictive power for stock returns in addition to dividend yields and interest rates;[17] however, we did not include a fourth state variable in our analysis on account of limitations in computing power. Denoting the rate of return on the stock portfolio by $\mathrm{d}S/S$, the joint stochastic process for the state variables is assumed to be of the form:

$$\frac{\mathrm{d}S}{S} = \mu_s \, \mathrm{d}t + \sigma_s \, \mathrm{d}z_S, \tag{3}$$

$$\mathrm{d}r = \mu_r \, \mathrm{d}t + \sigma_r \, \mathrm{d}z_r, \tag{4}$$

$$\mathrm{d}l = \mu_l \, \mathrm{d}t + \sigma_l \, \mathrm{d}z_l, \tag{5}$$

$$\mathrm{d}\delta = \mu_\delta \, \mathrm{d}t + \sigma_\delta \, \mathrm{d}z_\delta, \tag{6}$$

where the parameters μ_i, σ_i $(i = r, l, \delta, S)$ are at most functions of the state variables r, l, δ, and $\mathrm{d}z_\mathrm{i}$ are increments to Wiener processes. The correlation coefficients between the increments to the Wiener processes are denoted by $\rho_{r\delta}$, etc.

The three asset classes assumed to be available to the investor for investment are cash with sure rate of return, r; stock, whose rate of return is given by Eq. (3); and consol bonds. The price of a consol bond, $B(l)$, is inversely proportional to its yield, l. The total return on a consol bond is the sum of the yield and the price change; then a simple application of Ito's lemma implies that the instantaneous

[14] An early study drawing attention to the importance of this variable is Lintner (1975). More recent studies include Keim and Stambaugh (1986) and Hodrick (1991). Attempts to account for this empirical regularity include Geske and Roll (1983) and Fama (1981).

[15] See, for example, Fama and French (1988).

[16] See Brennan and Schwartz (1982); this is a natural implication of expectations-based theories of the term structure.

[17] See Keim and Stambaugh (1986).

total return on the consol bond is given by

$$\frac{\mathrm{d}B}{B} + l\,\mathrm{d}t = \left(l - \frac{\mu_l}{l} + \frac{\sigma_l^2}{l^2}\right)\mathrm{d}t - \frac{\sigma_l}{l}\,\mathrm{d}z_l. \tag{7}$$

Define x as the proportion of the investment portfolio that is invested in stock, and y the proportion that is invested in the consol bond. Then the stochastic process for wealth, W, is:

$$\begin{aligned}
\frac{\mathrm{d}W}{W} &= \left[x\frac{\mathrm{d}S}{S} + y\left(\frac{\mathrm{d}B}{B} + l\,\mathrm{d}t\right) + (1 - x - y)r\,\mathrm{d}t\right. \\
&= \left[x(\mu_s - r) + y\left(1 - r - \frac{\mu_l}{l} + \frac{\sigma_l^2}{l^2}\right) + r\right]\mathrm{d}t \\
&\quad + \left[x^2\sigma_s^2 + y^2\frac{\sigma_l^2}{l^2} - \frac{2xy\sigma_l\sigma_s\rho_{sl}}{l}\right]^{1/2}\mathrm{d}z_w \\
&\equiv \mu_w\,\mathrm{d}t + \sigma_w\mathrm{d}z_w.
\end{aligned} \tag{8}$$

Define $V(r, l, \delta, W, \tau)$ as the expected utility under the optimal policy when there are τ periods to the horizon. Then the Bellman equation is

$$\operatorname*{Max}_{x,\,y}\ \mathrm{E}[\mathrm{d}V] = 0 \tag{9}$$

or,

$$\begin{aligned}
\operatorname*{Max}_{x,\,y}\ &[V_w\mu_w W + V_r\mu_r + V_l\mu_l + V_\delta\mu_\delta - V_\tau \\
&+ \tfrac{1}{2}V_{ww}\sigma_w^2W^2 + \tfrac{1}{2}V_{ll}\sigma_r^2 + \tfrac{1}{2}V_{ll}\sigma_l^2 + \tfrac{1}{2}V_{\delta\delta}\sigma_\delta^2 \\
&+ V_{W\delta}W\sigma_\delta\sigma_w\rho_{\delta w} + V_{rl}\sigma_r\sigma_r\sigma_l\rho_{rl} + V_{r\delta}\sigma_r\sigma_\delta\rho_{r\delta} \\
&+ V_{Wr}W\sigma_r\sigma_w\rho_{rw} + V_{wl}W\sigma_l\sigma_w\rho_{lw} + V_{l\delta}\sigma_l\sigma_\delta\rho_{l\delta}] = 0.
\end{aligned} \tag{10}$$

As specified, the control problem (10) has four state variables including W. In order to reduce the number of state variables, we assume that utility is of the isoelastic form so that

$$V(r, l, \delta, W, 0) = \frac{1}{\gamma}W^\gamma, \quad \text{for } \gamma < 1. \tag{11}$$

Then it may be verified that $V(r, l, \delta, W, \tau)$ may be written as $\gamma^{-1}W^\gamma v(r, l, \delta, \tau)$, where

$$v(r, l, \delta, 0) = 1 \tag{12}$$

and

$$
\begin{aligned}
\operatorname*{Max}_{x,\,y}\Bigg[\, & \mu_w v + \frac{1}{\gamma}\mu_r v_r + \frac{1}{\gamma}\mu_l v_l + \frac{1}{\gamma}\mu_\delta v_\delta - \frac{1}{\gamma} v_\tau \\
& + \frac{1}{2}(\gamma - 1)\sigma_w^2 v + \frac{1}{2\gamma}\sigma_r^2 v_{rr} + \frac{1}{2\gamma}\sigma_l^2 v_{ll} + \frac{1}{2\gamma}\sigma_\delta^2 v_{\delta\delta} \\
& + \sigma_w \sigma_r \rho_{rw} v_r + \sigma_w \sigma_l \rho_{lw} v_l + \sigma_w \sigma_\delta \rho_{\delta w} v_\delta \\
& + \frac{1}{\gamma}\sigma_r \sigma_l \rho_{rl} v_{rl} + \frac{1}{\gamma}\sigma_r \sigma_\delta \rho_{r\delta} v_{r\delta} + \frac{1}{\gamma}\sigma_l \sigma_\delta \rho_{l\delta} v_{l\delta} \Bigg] = 0, \qquad (13)
\end{aligned}
$$

where

$$
\begin{aligned}
\sigma_w \sigma_r \rho_{rw} &\equiv \sigma_{wr} = x\sigma_s \sigma_r \rho_{rs} - y\frac{\sigma_l}{l}\sigma_r \rho_{rl}, \\
\sigma_w \sigma_l \rho_{rw} &\equiv \sigma_{wl} = x\sigma_s \sigma_l \rho_{ls} - y\frac{\sigma_l^2}{l}, \qquad (14) \\
\sigma_w \sigma_\delta \rho_{\delta w} &\equiv \sigma_{w\delta} = x\sigma_s \sigma_\delta \rho_{s\delta} - y\frac{\sigma_l}{l}\sigma_\delta \rho_{\delta l}.
\end{aligned}
$$

Substituting for μ_w and collecting terms, we have finally:

$$
\begin{aligned}
\operatorname*{Max}_{x,\,y}\Bigg\{ & v\Bigg[x(\mu_s - r) + y\left(l - r - \frac{\mu_l}{l} + \frac{\sigma_l^2}{l^2}\right) + r \\
& + \frac{1}{2}(\gamma - 1)\left(x^2\sigma_s^2 + y^2\frac{\sigma_l^2}{l^2} - \frac{2xy\sigma_s\sigma_l\rho_{sl}}{l}\right)\Bigg] \\
& + v_r\left[\frac{1}{\gamma}\mu_r + x\sigma_{sr} - \frac{y}{l}\sigma_{rl}\right] + v_l\left[\frac{1}{\gamma}\mu_l + x\sigma_{sl} - \frac{y}{l}\sigma_l^2\right] \\
& + v_\delta\left[\frac{1}{\gamma}\mu_\delta + x\sigma_{s\delta} - \frac{y}{l}\sigma_{l\delta}\right] - \frac{1}{\gamma}v_\tau \\
& + \frac{1}{\gamma}\left[\frac{1}{2}v_{rr}\sigma_r^2 + \frac{1}{2}v_{ll}\sigma_{ll}^2 + \frac{1}{2}v_{\delta\delta}\sigma_\delta^2 + v_{rl}\sigma_{rl} + v_{r\delta}\sigma_{r\delta} + v_{l\delta}\sigma_{l\delta}\right]\Bigg\} = 0. \qquad (15)
\end{aligned}
$$

The first-order conditions for a maximum in (15) imply that the optimal controls, $x^* \equiv x^*(r, l, \delta, \tau)$ and $y^* \equiv y^*(r, l, \delta, \tau)$, are given by

$$
\begin{aligned}
x^* = \frac{1}{(\gamma - 1)(\sigma_{sl}^2 - \sigma_s^2\sigma_l^2)}\Bigg[& \frac{\sigma_l^2}{l}(\mu_s - r) + \sigma_{sl}\left(l - r - \frac{\mu_l}{l} + \frac{\sigma_l^2}{l^2}\right) \\
& + \frac{v_r}{lv}(\sigma_{sr}\sigma_l^2 - \sigma_{rl}\sigma_{sl}) + \frac{v_\delta}{lv}(\sigma_{s\delta}\sigma_l^2 - \sigma_{ls}\sigma_{l\delta})\Bigg] \qquad (16)
\end{aligned}
$$

$$y^{+} = \frac{l^2}{(\gamma - 1)(\sigma_{sl}^2 - \sigma_s^2\sigma_l^2)}\left[\frac{\sigma_{sl}}{l}(\mu_s - r) + \sigma_s^2\left(l - r - \frac{\mu_l}{l} + \frac{\sigma_l^2}{l^2}\right)\right.$$
$$\left. + \frac{v_r}{lv}(\sigma_{sr}\sigma_{sl} - \sigma_{rl}\sigma_s^2) + \frac{v_\delta}{lv}(\sigma_{s\delta}\sigma_{sl} - \sigma_s^2\sigma_{ls}) + \frac{v_l}{lv}(\sigma_{sl}^2 - \sigma_l^2\sigma_s^2)\right]. \tag{17}$$

4. The stochastic process

In order to estimate the joint stochastic process (3)–(6) for the state variables and the stock return, it is necessary to specify the functional forms of the drift and diffusion coefficients. The basic assumption we made was that the expected returns on stocks and bonds, and the drifts of the dividend yield and short rate, were linear functions of the three state variables, r, l, and δ, while the volatility of each state variable was assumed to be proportional to its current level, and the volatility of the stock rate of return was taken as constant. This implies from Eq. (7) that the drift of the long rate is a non-linear function of the state variables, being equal to the product of l and a linear function of the state variables. This specification implies that the joint stochastic process may be written as

$$\frac{dS}{S} = (a_{s1} + a_{s2}\delta + a_{s3}r + a_{s4}l)\,dt + \sigma_s dz_s, \tag{18}$$

$$dr = (a_{r1} + a_{r2}\delta + a_{r3}r + a_{r4}l)\,dt + r\sigma_r dz_r, \tag{19}$$

$$dl = l(a_{l1} + a_{l2}\delta + a_{l3}r + a_{l4}l)\,dt + l\sigma_l dz_l, \tag{20}$$

$$d\delta = (a_{\delta 1} + a_{\delta 2}\delta + a_{\delta 3}r + a_{\delta 4}l)\,dt + \delta\sigma_\delta dz_\delta. \tag{21}$$

The dividend yield is defined as the sum of the past 12 months' dividends divided by the current level of the stock index, S. The specification (21) must therefore be regarded as an approximation since the stochastic process for lagged dividends is not modelled explicitly; to have done so would have introduced a fourth state variable into the analysis which would have considerably increased the difficulty of solving the control problem. However, we expect that the stochastic increment to the dividend yield will have a strong negative correlation with the return on the stock, since most of the stock return is accounted for by price changes.

The joint stochastic process was estimated by using a discrete approximation to the continuous process, and using monthly data for the period January 1972 to December 1991. The stock return was taken as the rate of return on CRSP *value weighted market index*. The short rate was taken as the yield on a one month Treasury Bill which was taken from the CRSP Government Bond File. The long rate was taken as the yield to maturity on the longest maturity taxable, non-callable government bond, excluding flower bonds; bond yield data were from the CRSP Government Bond File. The dividend yield was defined as the sum of the past 12 months' dividends on the CRSP *value weighted index*, divided by the current value of the index.

The system of Eqs. (18)–(21) was estimated by non-linear seemingly unrelated regression using TSP. Table 2 reports the regression estimates and Table 3 contains the estimated correlations of the innovations. As previous investigators have found, the expected return on common stocks is negatively related to the current level of the short rate and positively related to the level of the dividend yield, but is not significantly related to the long rate. As Brennan and Schwartz (1982) have found, the change in the short rate is negatively related to its current level and positively related to the level of the long rate – thus the short rate tends to adjust towards the long rate. The change in the long rate itself is the least predictable of our series, being negatively related to its current level and positively related to the short rate at marginal levels of significance. The change

Table 2
The estimated stochastic process for the state variables and the market return: January 1972–December 1991

		δ	r	l	σ
$\mathrm{d}S/S$	− 0.022	1.707	− 0.513	− 0.017	0.045
	(1.63)	(3.95)	(3.25)	(0.08)	
$\mathrm{d}r$	− 0.001	0.0048	− 0.216	0.181	0.134
	(0.34)	(0.05)	(− 5.66)	(3.54)	
$\mathrm{d}l$	0.025	− 0.318	0.247	− 0.308	0.037
	(1.82)	(0.79)	(1.68)	(1.55)	
$\mathrm{d}\delta$	0.0009	− 0.059	0.023	0.0005	0.048
	(1.62)	(2.97)	(3.27)	(0.05)	

t-statistics in parentheses.

Table 3
Correlations of state variable and stock return innovations: January 1972–December 1991

	Stock return	r	l	δ
Stock return	1.0			
r	− 0.037	1.0		
l	− 0.330	0.330	1.0	
δ	− 0.995	0.032	0.298	1.0

in the dividend yield is negatively related to its current level, so that it shows mean reversion; in addition, it is positively related to the short rate. As anticipated, the innovation in the dividend yield is very highly negatively correlated with the innovation in stock returns; it is also positively correlated with the innovation in the long rate, because the innovation in the long rate is negatively correlated with the innovation in stock returns. The remaining innovations have very low correlations.

In solving the control problem we shall truncate the state space by eliminating states in which the variables assume high values, beyond the range of US experience. It is important therefore that the state variables under the estimated stochastic process be stationary and not tend towards the boundary of our truncated state space if they are started at values within the range of historical experience. It is not possible to evaluate formally the stability of the stochastic differential equation system on account of the non-linearity entering through the equation for l. Therefore, we followed the empirical procedure of starting the system at points corresponding to historical joint realizations of the state variables, and then simulating the system forward while setting the innovations equal to zero: in all cases the system converged. Figs. 2 and 3 show the results of

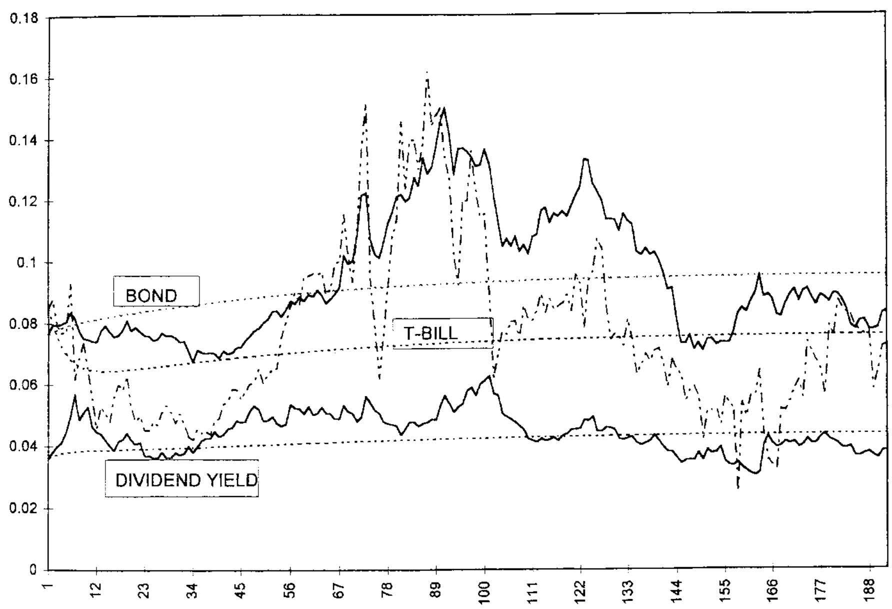

Fig. 2. The evolution of the state variables r, l and δ starting from their values on 29 March 1974 to December 1991, as determined by the nonstochastic part of the process (18)–(21) with parameter values as reported in Table 2. For comparison, the actual evolution of the state variables is also shown.

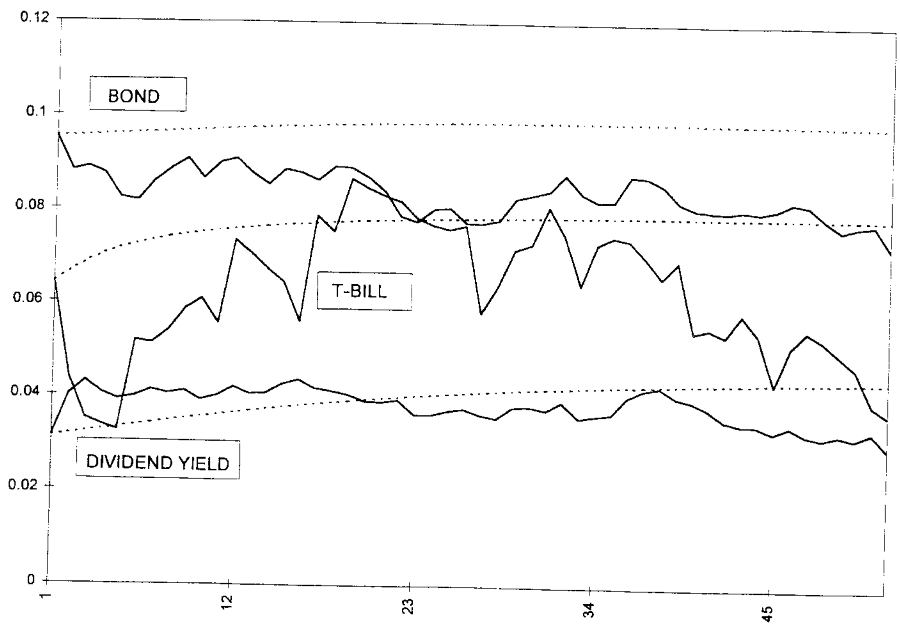

Fig. 3. The evolution of the state variables r, l and δ starting from their values on 30 September 1987 to December 1991, as determined by the nonstochastic part of the process (18)–(21) with parameter values as reported in Table 2. For comparison, the actual evolution of the state variables is also shown.

two representative simulations, with the actual realizations of the state variables being shown along with the simulated values. It can be seen that the system settles down to its long-run steady state which corresponds roughly to $l = 9\%$, $r = 8\%$, and $\delta = 4\%$. We conclude that the system is sufficiently well behaved to provide a useful input to our model. Fig. 4 shows the time series of annualized expected returns on the three asset classes that are implied by the model parameters. The expected returns on bonds and stock generally exceed the cash return, the notable exceptions being at the beginning of the 1970s when the expected returns on both bonds and stock are negative and, more particularly, at the beginning of the 1980s, when the very high short rate drives the estimated expected return on stocks negative for a prolonged period.[18]

In order to test the stability of the stochastic process, the system (18)–(21) was re-estimated for the two halves of the sample period, and the parameter

[18]Boudoukh et al. (1993) also report 'reliable evidence that the ex-ante risk premium is negative in some states of the world; these states are related to periods to high expected inflation and especially to downward-sloping term structures'.

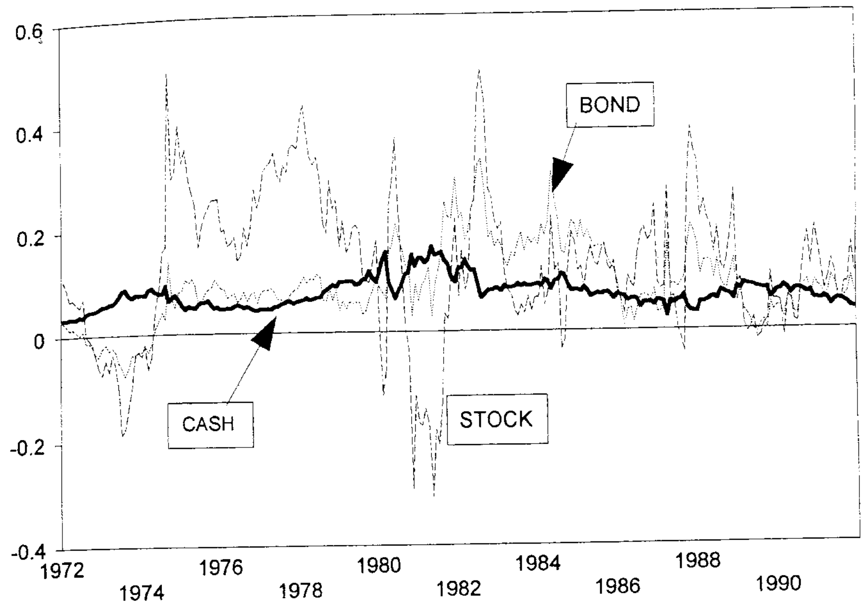

Fig. 4. For each month the annualized expected returns on bonds, stock and cash implied by the stochastic processes (18), (20), and the contemporaneous values of the state variables r, l and δ. The parameter values for the stochastic processes are reported in Table 2.

Table 4a
Estimates for the first half of the sample period: January 1972–December 1981

	Constant	δ	r	l	σ
dS/S	– 0.040	1.356	– 0.624	0.420	0.044
	(2.01)	(2.63)	(2.66)	(1.06)	
dr	0.000	– 0.090	– 0.067	0.108	0.106
	(0.08)	(1.00)	(1.29)	(1.18)	
dl/l	0.010	– 0.156	0.188	– 0.120	0.036
	(0.54)	(0.33)	(0.80)	(0.32)	
$d\delta$	0.001	– 0.033	0.027	– 0.018	0.048
	(1.50)	(1.42)	(2.56)	(0.91)	

Log-likelihood: 908.
t-ratios in parentheses.

estimates are reported in Table 4a and 4b. While the coefficients are generally similar across the two sub-periods, we note that both the effect of the dividend yield on the stock return and the effect of the level of the short rate on the change in itself are much greater in the second half of the sample period. A likelihood

Table 4b
Estimates for the second half of the sample period: January 1982–December 1991

	Constant	δ	r	l	σ
$\mathrm{d}S/S$	− 0.031	3.866	− 0.720	− 0.662	0.044
	(1.35)	(4.11)	(2.39)	(1.59)	
$\mathrm{d}r$	0.000	− 0.011	− 0.384	0.272	0.146
	(0.08)	(0.04)	(6.41)	(2.75)	
$\mathrm{d}l/l$	0.045	− 1.503	0.223	− 0.037	0.037
	(2.10)	(1.66)	(0.81)	(0.10)	
$\mathrm{d}\delta$	0.002	− 0.170	0.031	0.028	0.047
	(1.81)	(4.00)	(2.35)	(1.61)	

Log-likelihood: 828.
t-ratios in parentheses.

ratio test of the equality of the coefficients across the subperiods yields a χ^2 statistic of 58 with 16 degrees of freedom which is sufficient to reject the null hypothesis of constant coefficients at the 1% level. In an effort to allow for the effects of time variation in the coefficients we shall examine the performance of the optimal strategy when it is based on out-of-sample parameter estimates.

5. Numerical solution of the control problem

Substitution of the expressions (16) and (17) for the optimal controls into expression (15) yields a non-linear partial differential equation for v. The equation was solved using an implicit finite difference approximation on a $(40 \times 40 \times 20)$ grid. The second-order partial derivatives with respect to the state variables were discretized using second-order accurate central difference approximations. The first-order partial derivatives with respect to the state variables, on the other hand, were discretized using first-order accurate upwind difference approximations. This choice of upwind differencing was found to enhance the stability and convergence rate of the relaxation method used to solve the discrete system of equations. The interest rates were allowed to range from zero to 20% and the dividend yield from zero to 10%. Thus the step size in each of the state variables was 0.5%. The time step was set at approximately two weeks (1/24 yr). For each time step, initial trial values of the controls, x^* and y^*, were computed using values of the partial derivatives from the previous time step. Current values of $v(\delta, r, l, \tau)$ were then calculated from the partial differential equation using successive over relaxation, and the controls were then re-computed using values of the partial derivatives from the currently computed values of v ( ). The new controls were then used to compute new values of v( ); this procedure was repeated for a total of three iterations, and yielded satisfactory convergence.

In solving the equation, the following boundary conditions were imposed:

$$v_r = 0 \quad \text{at } r = 0,$$

$$v_{rr} = 0 \quad \text{at } r = 0.2,$$

$$v_l = 0 \quad \text{at } l = 0,$$

$$v_{ll} = 0 \quad \text{at } l = 0.2,$$

$$v_\delta = 0 \quad \text{at } \delta = 0,$$

$$v_{\delta\delta} = 0 \quad \text{at } \delta = 0.1.$$

For a 20 yr horizon (480 time steps) problem the program, which was written in FORTRAN, required about two CPU hours on a SUN SPARCstation 10 or about six hours on a Pentium PC.

6. Results

The control problem was solved for a value of γ equal to -5, both with and without constraints on short positions, and with a horizon of 20 yr. This rather extreme value of the risk aversion coefficient was chosen to offset our treatment of the parameters of the stochastic process as known when they are in fact estimated and therefore subject to estimation error.[19] The main results are shown in Figs. 5–7 which show the optimal portfolio proportions when no short sales are allowed, for the values of the state variables realized over the sample period.

The portfolio proportions are calculated for three distinct strategies. First, under the assumption that the horizon is a constant 20 years – the '20 year' strategy. Secondly, under the assumption that the horizon is always 1 month – the '1 month' strategy: this strategy is intended to represent the myopic strategy that underlies tactical asset allocation. Finally, under the assumption that the horizon is 1 January 1992 – the '1992' strategy; under this strategy the horizon used to calculate the portfolio proportions in any given month is the number of months remaining to January 1992.

[19] For an analysis of estimation error in (single period) portfolio problems see Bawa et al. (1979). For intertemporal problems with learning see Gennotte (1986).

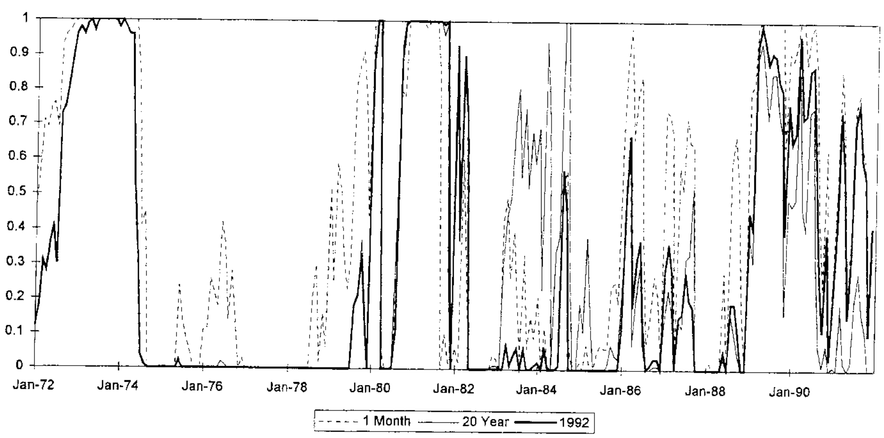

Fig. 5. For each month the proportion of the portfolio allocated to cash under the 1 month, 20 yr and 1992 strategies when the portfolio proportions are constrained to be non-negative. The parameter values for the stochastic processes used in the optimization are reported in Table 2.

Fig. 5 plots the cash proportions under the three strategies. We note that for all three strategies the cash ratio varies at least between zero and 90% and is highly volatile.[20] The one-month strategy usually involves a higher cash position than the 20 yr strategy. The reason for this is that cash is riskless over a one-month horizon, but it is not riskless for someone with a 20 yr horizon because of the uncertainty surrounding the re-investment rate.[21] As one would expect, the 1992 strategy cash proportion starts out identical with the 20 yr strategy and converges to that of the one-month strategy.

Fig. 6 plots the stock proportions which range between zero and 100% and again are highly volatile. The 20 yr strategy always invests more in stock than the one-month strategy, and the differences are often large; for example, in late 1979 the one-month strategy places only about 10% of the portfolio in stock while the 20 yr strategy places 65%. As before, the 1992 strategy is intermediate

---

[20]Our analysis does not include costs of transacting in the securities. However, as DuBois (1992) notes in connection with tactical asset allocation (TAA), the development of futures markets means that 'currently, transaction costs of TAA implementations appear to be less than a tenth of those that existed, say, 10 years ago. Hence a significant impediment to the viability of TAA strategies has been largely eliminated for today's TAA managers'. Clarke (1992) estimates the round-trip costs in futures markets at 7–12 basis points including market impact.

[21]See Hicks (1946) for the distinction between capital risk and income risk – the former is relevant for someone with a short horizon while the latter becomes more important as the horizon lengthens.

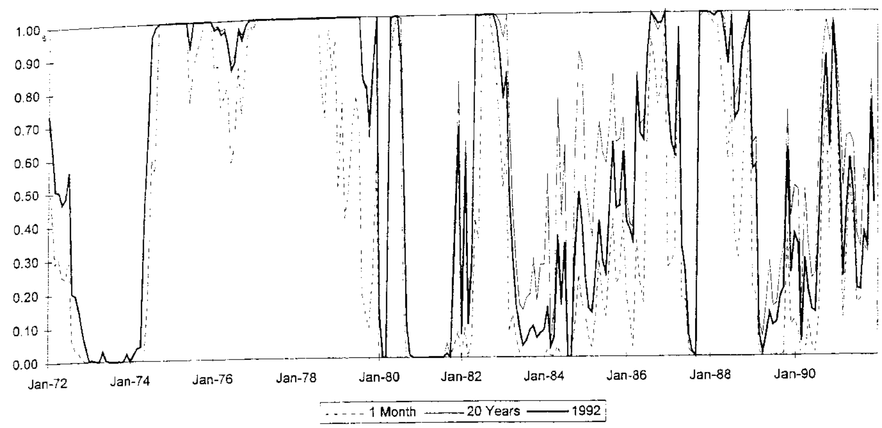

Fig. 6. For each month the proportion of the portfolio allocated to stock under the 1 month, 20 yr and 1992 strategies when the portfolio proportions are constrained to be non-negative. The parameter values for the stochastic processes used in the optimization are reported in Table 2.

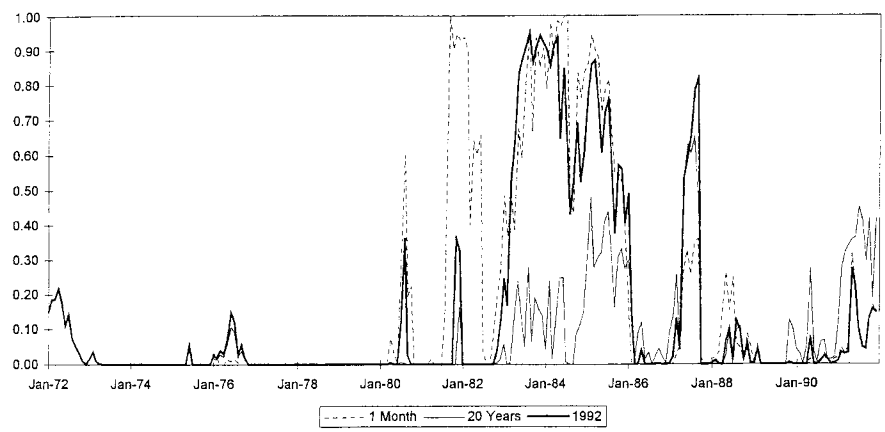

Fig. 7. For each month the proportion of the portfolio allocated to bonds under the 1 month, 20 yr and 1992 strategies when the portfolio proportions are constrained to be non-negative. The parameter values for the stochastic processes used in the optimization are reported in Table 2.

between these two. The intuition for the greater investment in stock for the long-horizon strategy is that the mean reversion in stock prices induced by the dividend yield variable means that the volatility of the distribution of (the log of) future (dividend adjusted) stock prices grows less than proportionately with time, so that stocks are less risky for those with a long horizon.

Fig. 7 plots the bond proportions. These are close to zero through the 1970s, then high and volatile in the 1980s, and low again in the 1990s. While intuition suggests that the 20 yr strategy would always assign a larger proportion to bonds than the one-month strategy, it is interesting to note that this is not always the case. For example, in January 1982 the 20 yr strategy assigns about 15% of the portfolio to bonds, while the one month strategy assigns around 90%, and the one-month strategy allocation to bonds exceeds that of the 20 yr strategy through most of the first half of the 1980s. The difference in the bond allocations is highly variable over time, suggesting that rules of thumb to adjust tactical asset allocation portfolios to account for a longer horizon are unlikely to be successful.[22]

In order to assess the performance of the strategy when out-of sample parameter estimates for the stochastic process are used to derive the optimal strategy, the optimal strategy was computed for each year of the second half of the sample period, January 1982 – December 1991, using parameter values estimated over the previous ten years. Thus the parameter estimates on which the optimal strategy is based are updated annually, always using the data on dividend yield and interest rates from the previous ten years. This strategy was then combined with the current values of the state variables to compute the optimal portfolio proportions for each month of the second half of the sample period, and the returns on this 'out-of-sample' strategy were calculated. The average monthly turnover was 89.6% per month for the unconstrained strategy, but only 22.5% per month for the constrained strategy. Figs. 8 and 9 report the out-of sample wealth relatives for the 1-month, 10 yr, and 1992 strategies along with the wealth relative for a pure stock investment for comparison. Fig. 9 relates to a strategy in which the portfolio proportions are constrained to be non-negative, while Fig. 8 relates to an unconstrained strategy. It may be seen in Fig. 8 that the unconstrained 1992 strategy performs slightly worse than the one-month strategy, but better than the other strategies over this sample period. However, when the strategies are constrained from taking negative positions, the 1992 strategy yields the highest final wealth outcome, as shown in Figure 9. The standard deviations and means of the monthly returns of the different strategies are reported in Table 5. It may be seen first that the standard deviation of the 1992 and 10 yr strategies are larger than those of the one-month strategy; as mentioned previously, the long-term strategic strategies accept more short-term risk than the one-month strategy because of the mean reversion in stock prices and relation between negative bond returns and higher future interest rates. The constrained strategies have both lower means and lower standard

[22]Some practitioners of tactical asset allocation (TAA) for pension funds construct a duration matched bond portfolio to mimic the liability portfolio, and practice TAA with the portfolio surplus. The results here suggest that this will not approximate the optimal strategy for a long horizon.

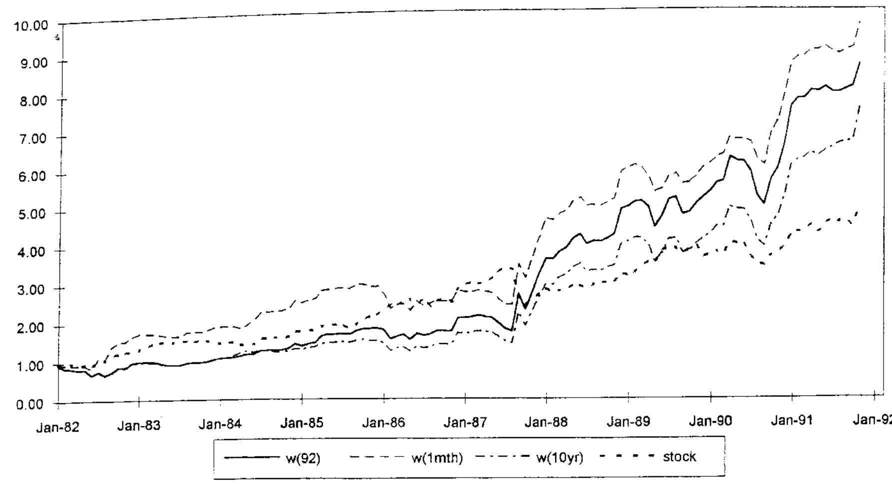

Fig. 8. The wealth relatives obtained by following the optimal 1992 strategy ($w(92)$), the optimal one-month strategy ($w(1\mathrm{mth})$), the optimal constant 10 yr horizon strategy ($w(10\mathrm{yr})$) and an all stock strategy (stock). The optimal strategies for each calendar year are determined using parameter estimates for the stochastic process (18)–(21) obtained using 10 yrs' data ending the previous December. The portfolio proportions are not constrained to be non-negative.

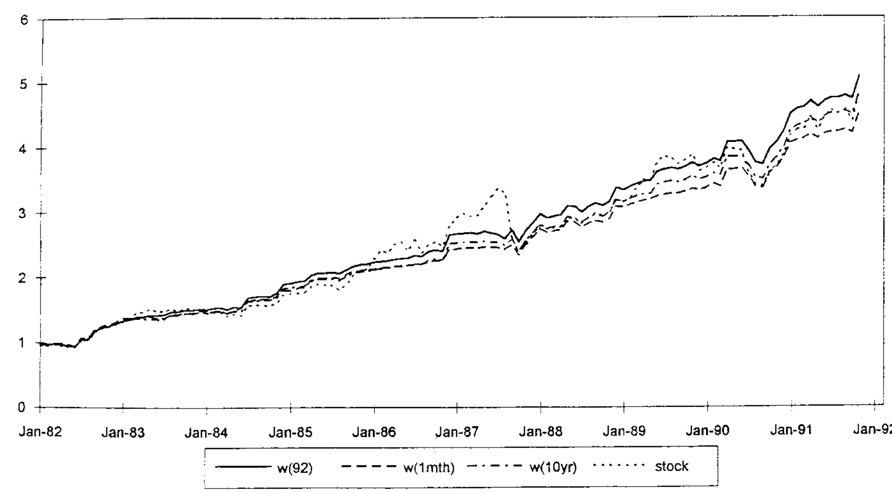

Fig. 9. The wealth relatives obtained by following the optimal 1992 strategy ($w(92)$), the optimal 1 month strategy ($w(1\mathrm{mth})$), the optimal constant 10 yr horizon strategy ($w(10\mathrm{yr})$) and an all stock strategy (stock). The optimal strategies for each calendar year are determined using parameter estimates for the stochastic process (18)–(21) obtained using 10 yrs' data ending the previous December. The portfolio proportions are constrained to be non-negative.

Table 5
Means (μ) and standard deviations (σ) of monthly returns under alternative out of sample investment strategies 1982–1991

Strategy	σ	μ
Cash	0.001	0.005
Bonds	0.035	0.012
Stock	0.047	0.015
Unconstrained: one-month horizon	0.071	0.022
Unconstrained: 1992 horison	0.083	0.022
Unconstrained: 10 yr horizon	0.084	0.021
Constrained: one-month horizon	0.028	0.013
Constrained: 1992 horizon	0.031	0.015
Constrained: 10 yr horizon	0.031	0.014

deviations than the unconstrained strategies. Some indication of the out of sample power of the model is provided by the fact that the constrained 1992 strategy has the same mean return as the all stocks strategy, but a monthly standard deviation of 3.1% as compared with 4.7% for the all stock strategy.

However, strictly speaking, it is inappropriate to compare the different strategies on the basis of a mean-variance criterion since, while the one-month strategy is approximately mean variance efficient,[23] the long-term strategies maximize the expected value of a (derived) utility function that depends on the state variables as well as on wealth. To see the importance of the distinction note that a strategy that invested in a pure discount bond which matured at the horizon would be riskless, but the standard deviation of the monthly returns would not be zero if there was variation in the interest rate. Therefore a better metric by which to assess the relative risk of the strategies is the variance of the *certainty equivalent* of wealth, $CE(W, r, l, \delta, \tau)$. The certainty equivalent is that amount of wealth such that the investor is indifferent between receiving it for sure at the horizon, and having his current wealth today *and* the opportunity to invest it optimally up to the horizon. Thus, the certainty equivalent is defined by

$$\frac{1}{\gamma}(CE)^{\gamma} = V(W, r, l, \delta, \tau) = \frac{1}{\gamma} W^{\gamma} v(r, l, \delta, \tau).$$

Hence

$$CE(W, r, l, \delta, \tau) = W\,[v(r, l, \delta, \tau)]^{1/\gamma}.$$

[23]With a short decision horizon, the objective function may be closely approximated by a function which is quadratic in wealth and does not depend on the other state variables.

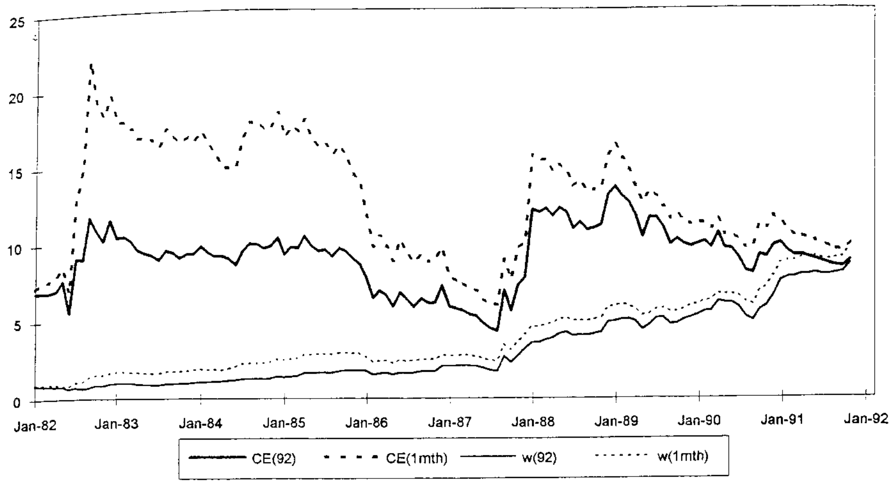

Fig. 10. The certainty equivalent wealth (CE) and realized wealth (w) under the optimal 1992 strategy (92) and the optimal 1 month strategy (1mth) when the portfolio proportions are not constrained to be non-negative. The certainty equivalent is such that the investor is indifferent between receiving it for sure at the horizon and having current wealth and the opportunity to invest in the available opportunities up to the horizon.

Fig. 10 plots the certainty equivalents for the 1992 and the one-month unconstrained strategies, where the latter is calculated by multiplying the wealth realized under the one-month strategy by $v^{1/\gamma}$, and v, which is computed from the 1992 strategy, takes account of the value of the investment opportunities remaining till 1992. Fig. 10 shows strikingly that in the early years the volatility of the certainty equivalent is much greater than the volatility of wealth: in other words, most of the risk in the early years comes, not from changes in wealth, but from changes in future investment prospects as captured by the state variables. As the horizon is approached, variability in the value of the investment opportunity set is correspondingly reduced. Moreover, it is clear by inspection that the volatility of the certainty equivalent under the one-month strategy, which does not attempt to hedge against shifts in the investment opportunity set, far exceeds the volatility of the certainty equivalent under the 1992 strategy. Thus, the 1992 strategy, as expected, does a better job of hedging against changes in the investment opportunity set.

Fig. 11 plots CE/W, the certainty equivalent per dollar of wealth for the constrained (ce^*) and unconstrained strategies (ce), along with the state variables. The normalized certainty equivalent, ce, is a measure of the value of the

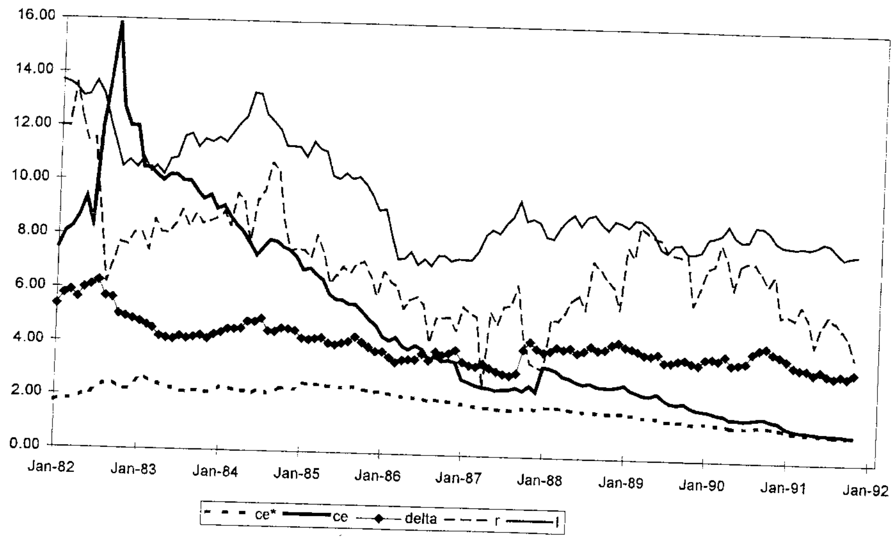

Fig. 11. The normalized certainty equivalent wealth ($ce = CE/w$) under the optimal 1992 strategy when the portfolio proportions are constrained to be non-negative (ce^*) and are not constrained (ce). The certainty equivalent is such that the investor is indifferent between receiving it for sure at the horizon and having one dollar of wealth and the opportunity to invest in the available opportunities up to the horizon. Also shown are the values of the state variables r, l and δ.

future investment opportunities for an investor with a given horizon and utility function.[24] Notice that in the early years the normalized certainty equivalent is much greater for the unconstrained strategy, reflecting its ability to take more advantage of investment opportunities by taking short positions; it is also more sensitive to changes in the state variables. We also observe that the normalized certainty equivalent tends to decline with time, reflecting the reduction in future investment opportunities, eventually reaching unity at the horizon.

An investment strategy that hedges perfectly against changes in the investment opportunity set will ensure that CE is constant. Since $CE = ce\,W$, this would imply that the change in the log of ce was equal and opposite in sign to the change in the log of W. To compare the hedging characteristics of the unconstrained 1992 and one-month strategies the monthly change in the logarithm of ce was regressed against the monthly change in the logarithm of wealth

---

[24]Note that the iso-elastic utility function ensures that CE/W is independent of W.

under the two strategies with the following results:[25]

$$\begin{aligned}
\operatorname{dln} ce &= \underset{(2.89)}{-\,0.008} + \underset{(1.40)}{0.119}\,\operatorname{dln} W_{1992}, \quad R^2 = 0.02,\\
\operatorname{dln} ce &= \underset{(3.82)}{-\,0.010} + \underset{(3.95)}{0.364}\,\operatorname{dln} W_{1\text{mth}}, \quad R^2 = 0.12.
\end{aligned}$$

We note that while the change in $\ln ce$ is positively related to the change in $\ln W$ under the one-month strategy, the relation is much less strong and is statistically insignificant under the 1992 strategy: this is consistent with the greater weight placed on hedging considerations under the 1992 strategy.

Fig. 11 clearly shows that the value of future investment opportunities is significantly influenced by variation in the state variables that we have chosen to
2 capture investment opportunities; however, it is difficult to discern from the figure the relative importance of the three state variables. To assess this, the logarithm of the estimated value of ce for the constrained strategy was regressed on the logarithms of the state variables and time to the horizon for each month from January 1982 to December 1991:

$$\ln ce = \underset{(2.24)}{-\,0.578} - \underset{(1.49)}{0.152}\ln\delta - \underset{(3.53)}{0.157}\ln r + \underset{(3.23)}{0.301}\ln l + \underset{(21.56)}{0.257}\ln\tau, \qquad R^2 = 0.87$$

ce is most strongly affected by the remaining time to maturity. It is increasing in l, which is a direct measure of the favorableness of investment opportunities. However, it is decreasing in r; this appears to be related to the fact that the expected return on stock is negatively related to the short-term interest rate. On the other hand, the dividend yield, which is positively related to the expected return on stock does not enter the regression significantly. It should be noted that the log–log specification of the regression is arbitrary.

While the results we have reported relate only to the single sample period, 1982–1991, they provide encouraging evidence that asset allocation strategies designed to take account of time variation in expected returns can provide significant performance improvement over static strategies, and comparison of the 1992 strategy with the myopic one-month strategy points to the importance of taking account of the investor's time horizon in devising optimal portfolio strategies.

---

[25] t-Ratios in parenthesis.

7. Conclusion

In this paper we have shown how it is possible to apply dynamic portfolio theory to the design of optimal portfolio strategies for an investor with a long-term horizon when there is time variation in the expected returns on different asset classes. We find that the investor's time horizon has a significant effect on the composition of the optimal portfolio. An investor with a long horizon typically places a larger fraction of the portfolio in both stocks and bonds than does a myopic investor. The reason for this is the mean reversion in both bond and stock returns that makes these assets less risky from the viewpoint of a long-term investor. Equivalently, investments in stocks and, more particularly, bonds provide the long-term investor with a hedge against future adverse shifts in the investment opportunity set – by buying long-term bonds the investor protects himself against declines in future interest rate. Myopic strategies such as those commonly employed in simple tactical asset allocation implementations neglect this role of long-term assets and misleadingly treat cash as a riskless asset; in reality, cash is riskless only for an investor with a one-period horizon.[26]

The out of sample simulation results provide encouraging evidence that the predictability of asset returns is sufficient for strategies that take it into account to yield significant improvements in portfolio returns. In this paper we have derived the optimal strategy assuming that the parameters of the return generating process are known rather than estimated. The next challenge is to extend the scope of the analysis to take account of estimation risk. We anticipate that this will reduce the tendency of the unconstrained model to take highly levered portfolio positions and will reduce the implied portfolio turnover rates. Our analysis takes account of only three asset classes. Extension to additional asset classes is straightforward so long as the expected returns on these asset classes can be expressed in terms of the same set of state variables. Extending the analysis to incorporate additional state variables is straightforward in principle, but significantly increases computational requirements.

References

Arnott, R.D., Bernstein, P.L., 1992. The role of liabilities: defining and managing pension fund risk. In: Arnott, R.D., Fabozzi F.J. (Eds.) Active Asset Allocation, Probus, Chicago.

---

[26] More sophisticated tactical asset allocation models applied to pension fund management (cf. Arnott and Bernstein, 1992) introduce a 'liability asset' to take account of the fixed nature of the obligations of a defined benefit pension plan, and then apply standard (myopic) portfolio theory to the pension fund surplus (defined as the difference between the present values of assets and liabilities), in addition to holding the liability asset which is typically a duration matched portfolio of bonds. The net effect is to increase the value of the bonds held in the portfolio.

Arrow, K.J., 1971. Essays in the Theory of Risk Bearing. North-Holland, Amsterdam.

Bekaert, G., Hodrick, R.J., 1992. Characterizing predictable components in excess returns on equity and foreign exchange markets. Journal of Finance 47, 467–509.

Boudoukh, J., Richardson, M., Smith, T., 1993. Is the ex-ante risk premium always positive?', Journal of Financial Economics, 34, 387–408.

Brennan, M.J., Schwartz, E.S., 1982. An equilibrium model of bond pricing and a test of market efficiency. Journal of Financial and Quantitative Analysis 17, 301–329.

Bawa, V.S., Brown, S.J., Klein, A., 1979. Estimation Risk and Optimal Portfolio Choice. North-Holland, Amsterdam.

Campbell, J.Y., Shiller, R.J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies, 1, 195–228.

Carino, D.R., Kent, T., Myers, D.H.,Stacey, C., Watanabe, K., Ziemba, W.T., 1994. Russell–Yasuda Kasai model: an asset-liability model for a Japanese insurance company using multi-stage stochastic programming. Interfaces, 24, 24–49.

Clarke, R.G., 1992. Asset allocation using futures markets. In: Arnott, A.D., Fabozzi F.J. (Eds.) Active Asset Allocation. Probus, Chicago.

Cochrane, J., 1990. Explaining the variance of price-dividend ratios, Working Paper, University of Chicago.

DuBois, C.H., 1992. Tactical asset allocation. In: Arnott, A.D., Fabozzi, F.J. (Eds.), in Active Asset Allocation, Probus, Chicago.

Fama, E.F., 1981. Stock returns, real activity, inflation and money. American Economic Review, 71, 545–565.

Fama, E.F., French, K.R., 1988. Dividend yields and expected stock returns. Journal of Financial Economics, 22, 3–26.

Gennotte, G., 1986. Optimal portfolio choice under incomplete information. Journal of Finance, 41, 733–749.

Geske, R., Roll, R., 1983, The fiscal and monetary linkage between stock returns and inflation. Journal of Finance, 38, 1–33.

Hicks, J.R., 1946. Value and Capital. Oxford University Press.

Jessup, E.R., Yang, D., Zenios, S.A. 1994. Parallel factorization of structured matrices arising in stochastic programming. SIAM Journal of Optimization, 4, 833–846.

Keim, D.B., and Stambaugh, R.F., 1986. Predicting returns in the stock and bond markets. Journal of Financial Economics, 17, 357–390.

Lintner, J., 1975. Inflation and security returns. Journal of Finance. 30, 259–280.

Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous time model. Journal of Economic Theory 3, 373–413.

Merton, R.C., 1990. Continuous Time Finance, Chapters 4–6. Basil Blackwell, Cambridge, MA.

Mulvey J.M., 1994. An asset-liability investment system. Interfaces, 24, 22–33.

Mulvey, J.M., A. Ruszczynski, 1995. A new scenario decomposition method for large-scale stochastic optimization. Operations Research. 43, 477–490.

Mulvey, J.M., Vladimirou, H., 1992. Stochastic network programming for financial planning problems. Management Science 38, 1642–1664.

Zenios, S.A., 1993. A model for portfolio management with mortgage-backed securities. Annals of Operations Research, 43, 337–356.